

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

MONGODB, INC.,)	
)	
Plaintiff,)	
)	
v.)	C.A. No.
)	
FERRETDDB INC.,)	
)	
Defendant.)	
)	

COMPLAINT

Plaintiff MongoDB, Inc. (“MongoDB” or “Plaintiff”) files this Complaint against FerretDB Inc. (“FerretDB” or “Defendant”) alleging as follows:

NATURE OF THE ACTION

1. For eighteen years, MongoDB has proudly developed and maintained an industry leading free and source-available database platform that is trusted by organizations around the world. MongoDB’s mission is to unleash the power of software and data, enabling software development teams to meet the diverse needs of modern applications.

2. MongoDB is a leading innovator in the field of digital databases. As is described in this Complaint, MongoDB pioneered the development of “non-relational” databases, which excel in scenarios where more conventional “relational” databases struggle, as in when dealing with the ever more diverse and massive datasets that are increasingly associated with the modern digital age.

3. As a result of MongoDB’s significant innovations—and the over \$1 billion it has invested to produce those innovations—MongoDB’s software has been downloaded hundreds of millions of times and is used by many millions of users. Likewise, MongoDB has tens of

thousands of paying customers, and it has partnerships with eminent technology companies around the world.

4. Since its inception, MongoDB has been committed to the benefits of having freely available source code for software development, including providing the software developer community with dependable access to its software and services. MongoDB believes that this approach fosters technological innovation and growth, both for MongoDB but even more so for the broader developer community. MongoDB offers a licensing arrangement to the developer community that is designed to promote the benefits of the source-available model by encouraging open collaboration, while prohibiting freeriders from selfishly exploiting MongoDB's technology for their own commercial gain.

5. Unfortunately, Defendant FerretDB has taken unfair, unlawful advantage of MongoDB's open approach and has crossed lines that MongoDB cannot ignore. For one, rather than breaking into the database landscape with its own competitive product, FerretDB has infringed on MongoDB's patented technology and declined to end its infringement even after MongoDB asked FerretDB to do so numerous times. FerretDB's patent infringement not only harms MongoDB's ability to deliver innovative products to its customers and community, but also impairs the trust and spirit of collaboration that are essential to the developer community.

6. Moreover, FerretDB has—since its launch in 2021—made harmful and false statements about its own products, MongoDB products, and the relationship between them, including, for example, falsely representing to the public that it has built a “replacement” product for MongoDB and falsely claiming that FerretDB is compatible with MongoDB in ways that it is not. In truth, FerretDB has built consistently inferior products to MongoDB's and has done so only by misappropriating MongoDB's technology.

7. Compounding the injury to MongoDB, FerretDB has improperly offered its subpar products by tarnishing and diluting MongoDB's brand, including its trademarks. Indeed, FerretDB's co-founder and CEO, Peter Farkas, has publicly stated that "one of the reasons [FerretDB] chose" the company's original product name, *MangoDB*, "is because, *obviously, MongoDB comes into mind.*"¹ Through this and other means described in this Complaint, FerretDB has also taken deliberate steps to interfere with MongoDB's ability to generate the revenue it needs to support the innovations that its customers, its users, and the developer ecosystem depend on.

8. MongoDB respects and supports legitimate competition, understanding that a vibrant free and source-available ecosystem depends on the ability of developers and companies to build upon one another's ideas, so long as they do so within the bounds of the law and the norms of community respect.

9. However, FerretDB has acted well beyond these bounds and norms. FerretDB's conduct strikes at the heart of free and source-available development itself. Source-available and open source development thrives when participants innovate, respect the rights of others, and clearly communicate with the community. FerretDB's disregard for these principles—by misleading FerretDB's and MongoDB's users and customers, diluting the hard-earned value of MongoDB's respected trademarks, and more—threatens the delicate ecosystem that has made successful projects like MongoDB possible.

10. MongoDB has attempted to reach a resolution with FerretDB for years, but FerretDB has largely ignored MongoDB, and has not only continued, but escalated, its unlawful

¹ Document Database Community, The Current State of MongoDB Alternatives, and 2Y of FerretDB, YouTube at 36:10 (Sept. 19, 2023), https://www.youtube.com/watch?v=43F1I9I6_gA.

conduct. As a result, MongoDB is forced to bring this action for patent infringement, trademark dilution, and false advertising.

THE PARTIES

11. Plaintiff MongoDB is a Delaware corporation with its principal place of business at 1633 Broadway, 38th Floor, New York, New York 10019. Founded in 2007, MongoDB offers the world's most popular document database.

12. Defendant FerretDB (formerly known as "MangoDB"), is a Delaware corporation and, upon information and belief, has a place of business located at 2055 Limestone Road, Suite 200C, Wilmington, Delaware 19808. FerretDB may be served with process through its registered agent, Northwest Registered Agent Service, Inc., at 8 The Green, Suite B, Dover, Delaware 19901.

JURISDICTION AND VENUE

13. Jurisdiction as to the patent infringement claims asserted herein is conferred on this Court by 28 U.S.C. §§ 1331 and 28 U.S.C. § 1338(a), respectively, and as to all other claims asserted herein pursuant to 28 U.S.C. § 1367(a).

14. This district has general and specific personal jurisdiction over FerretDB because FerretDB is incorporated under the laws of the State of Delaware, has committed acts within this district giving rise to this action, transacts and conducts business in this district and the State of Delaware, and transacts and conducts business with residents of this district and the State of Delaware.

15. Venue is proper in the District of Delaware pursuant to 28 U.S.C. § 1391(b)(1), and 28 U.S.C. § 1400(b) because FerretDB resides and maintains a place of business within the District of Delaware.

FACTUAL ALLEGATIONS

MongoDB's History of Innovation

16. MongoDB (formerly 10gen) was founded in 2007 by Dwight Merriman, Eliot Horowitz, and Kevin Ryan, veteran leaders of the pioneering online advertising company DoubleClick. While employed at DoubleClick, Merriman, Horowitz, and Ryan quickly learned the limitations of conventional “relational” databases as they worked to scale DoubleClick’s ad-serving technology to serve more than 400,000 ads to internet users every second.

17. Conventional relational databases were built for a world of relatively stable, lower-volume structured data. This world increasingly disappeared in the 2000s, as the modern digital ecosystem—with real-time applications, ever-increasing volumes of data, distributed users, and ever-evolving data formats—came to fruition. This newer evolution of the technology ecosystem strained the inherent limitations of conventional relational databases.

18. Enter MongoDB. After leaving DoubleClick, Merriman, Horowitz, and Ryan founded MongoDB in 2007 to pioneer a new database approach, using *non-relational* databases (sometimes known as “noSQL” databases). MongoDB launched its eponymous non-relational document database in 2009 with the goal of creating a next-generation data platform that could process large volumes of unstructured data, easily scale, and be flexible enough to meet the demands of modern applications.

19. Built by developers for developers, MongoDB is now one of the largest, publicly-held database companies in the world. The MongoDB database has been downloaded hundreds of millions of times, with millions of developers enrolling in MongoDB University. It is unsurprising, then, that industry analysts rank MongoDB as a category leader. MongoDB has forged partnerships with all the major cloud and technology companies around the world.

20. The most sophisticated organizations, from the world's largest companies to cutting-edge startups, use MongoDB every day to transform their companies, innovate new experiences, and disrupt their industries. Developers continue to demand MongoDB as one of the most loved and desired must-have technology platforms in their toolkit.

21. MongoDB's success has not come easily. It has invested more than \$1 billion in research and development and, as discussed below, millions of dollars more in getting the word out about MongoDB through its marketing and advertising.

FerretDB Infringes MongoDB's Patents

22. By offering an unlicensed purported "replacement" for MongoDB, FerretDB is infringing MongoDB's patent rights. As discussed above, MongoDB is a key pioneer in the field of distributed, document, and NoSQL database technology. These technologies power a wide variety of applications and platforms used in the development of web applications, APIs (which allow software systems to communicate with one another), and more. MongoDB spent many years developing its groundbreaking improvements to database design, management, and implementation, which resulted in a number of patents, including United States Patent Numbers 8,996,463 (the "'463 Patent"), 9,262,462 (the "'462 Patent"), 10,031,956 (the "'956 Patent"), and 10,866,868 (the "'868 Patent") (collectively, the "Asserted Patents").

23. As further described later in this Complaint, FerretDB is infringing all the Asserted Patents, which are directed to inventive aggregation framework system architectures and methods to optimize aggregate operations over a database based on data dependencies and other groundbreaking improvements to database design, management, and implementation.

24. When a company willfully and illegally misappropriates another company's technology to offer a purported "replacement" product, as FerretDB is doing, it directly harms

the innovator by depriving it of the revenues that fund its investment of time, creativity, and resources. But beyond the immediate damage to MongoDB, FerretDB’s conduct also hurts the broader developer community by violating the prevailing norms of collaboration and weakening respect for intellectual property rights, which are essential for fostering innovation. If companies were permitted to routinely copy others’ technology without regard for licenses or patents, it would discourage individuals and smaller developers from contributing new ideas—particularly in source-available and open source projects where creators often rely on a balance of openness and enforceable rights—ultimately stifling the healthy exchange of innovation that the broader technology ecosystem depends on. FerretDB’s infringement of MongoDB’s patents crosses a line that threatens to jeopardize this delicate balance.

FerretDB Misleadingly Represents Itself and MongoDB

25. FerretDB was founded in 2021 with two guiding principles. First, FerretDB sought to create a product that copied all of MongoDB’s technology. Second, it hoped to poach existing MongoDB customers by falsely representing that it had actually achieved that end, when in reality, FerretDB failed miserably.

26. FerretDB carries out these goals not only by infringing MongoDB’s patents, but also by falsely representing FerretDB’s functionality, misrepresenting MongoDB’s functionality, and falsely describing the relationship between these products.

27. FerretDB telegraphed its copycat goals as soon as it launched in 2021, stating that it sought “to match the behavior of MongoDB as closely as possible.”² But FerretDB has not in fact achieved its stated goal of replicating all of MongoDB’s functionality (much less the

² Peter Farkas, *FerretDB v.0.1.1 is out, and other exciting updates* (Apr. 18, 2022), <https://blog.ferretdb.io/ferretdb-v-0-1-1-is-out-and-other-exciting-updates/>.

quality of MongoDB’s offerings). That has not stopped FerretDB from falsely representing otherwise, however.

28. For example, today, FerretDB represents on its website that it “mimics the behavior of MongoDB *in every possible way*”³ and that FerretDB allows “the user [to] interact with FerretDB *just like* as if it would be MongoDB.”⁴ These representations by FerretDB are false and misleading for a number of reasons.

29. For one, and most fundamentally, MongoDB is a fully functional and revolutionary document database. It has all the components a developer needs. By contrast, FerretDB is a pale shadow of the fully-functional products offered by MongoDB. On its own, FerretDB is incapable of doing much of anything of value for developers.

30. FerretDB’s representations that it “mimics the behavior of MongoDB *in every possible way*”⁵ and allows “the user [to] interact with FerretDB *just like* as if it would be MongoDB”⁶ are false because—as shown by FerretDB’s own publicly available compatibility testing documentation—FerretDB does not offer all the tools, commands, and/or drivers provided by MongoDB. This documentation shows instead, for example, that even the latest release of FerretDB is *missing nearly a third* of the commands and tools supported by MongoDB’s database. For example, FerretDB provides *none* of the role management commands provided by MongoDB, including “createRole,” “dropAllRolesFromDatabase,” and many more, which FerretDB admits it has not implemented.⁷ For the same reasons, FerretDB’s representations on its website that it allows customers to “use *all* [the] tools and drivers [they]

³ <https://www.ferretdb.com/> (last visited Apr. 24, 2025) (emphasis added).

⁴ <https://www.ferretdb.com/about> (last visited Apr. 18, 2025) (emphasis added).

⁵ <https://www.ferretdb.com/> (last visited Apr. 24, 2025) (emphasis added).

⁶ <https://www.ferretdb.com/about> (last visited Apr. 18, 2025) (emphasis added).

⁷ See <https://docs.ferretdb.io/migration/compatibility/> (last visited Apr. 26, 2025).

would with MongoDB”⁸ and “the *same commands*, drivers, and tools as those of MongoDB”⁹ are also false.

31. Commands and tools are materially important features to the developer community. They are among the essential ways that developers interact with products like MongoDB’s and FerretDB’s, and the availability and breadth of them drives developers’ product-use decisions. Role management commands, for example, are crucial for managing user privileges and access control within a database. In other words, they define what MongoDB’s (and FerretDB’s) users can and cannot do. FerretDB’s false suggestions that it has all the same commands as MongoDB are materially misleading to the developer community, including MongoDB’s customers, FerretDB customers, and prospective users of both companies’ products.

32. FerretDB’s representation that “the user” can “interact with FerretDB just like” MongoDB is false and/or misleading for the additional reason that FerretDB v.2.0 is not compatible with the many machines that use processors sold by Arm, Inc. This affects FerretDB v.2.0’s compatibility with popular, important hardware that “the user” may use, including Apple’s modern generation of MacBooks and the Graviton Processor offered by Amazon Web Services. Users who rely on either system are unable to “interact with FerretDB just like as if it would be MongoDB,” as FerretDB falsely claims.

33. FerretDB makes yet more false representations. It touts that it “is designed to be a *drop-in replacement* for MongoDB, so you can keep using your existing tools and libraries *without any changes*.”¹⁰ And, likewise, it claims that FerretDB “is compatible with MongoDB

⁸ <https://www.ferretdb.com/> (last visited Apr. 24, 2025) (emphasis added).

⁹ *Id.* (emphasis added).

¹⁰ Alexander Fashakin, *Migrating MongoDB Data to FerretDB with dsync* (Apr. 16, 2025), <https://blog.ferretdb.io/migrate-mongodb-data-ferretdb-dsync/> (emphasis added).

drivers and can be *used as a direct replacement for MongoDB 5.0+.*”¹¹ These statements are false because, for the reasons discussed above, FerretDB fails to offer the commands and compatibility features that MongoDB offers. Therefore, FerretDB simply is not a “drop-in replacement” for MongoDB, nor can users use all “existing tools” that MongoDB offers, as FerretDB misleadingly claims. These misstatements—like the others discussed above—are material because developers who believe FerretDB’s false statement may choose to use its product, on the basis of false pretenses, instead of MongoDB’s offerings (or those of another company).

34. Moreover, FerretDB states that it allows users to “turn their Postgres services into MongoDB as a service.”¹² This is false for at least two reasons. First, FerretDB, of course, does *not* “turn Postgres”—a completely separate database product that MongoDB does not offer—into “MongoDB’s” product offering. Second, FerretDB—while representing that it intends to offer a cloud “service” in the *future*—does not in fact currently offer a product as a “service.” In the software industry, developers understand that when a software product is offered “as a service,” that generally means the software is offered remotely through the cloud, and typically on a subscription, recurring-payment basis, instead of being installed on local devices or server hardware, with less frequent payments for the software. Contrary to its false representation, FerretDB does not currently offer such a model.

35. FerretDB’s misrepresentations about certain core functionality of its products, and those of MongoDB’s, are harmful to MongoDB, to its customers, and to members of the

¹¹ <https://docs.ferretdb.io/> (last visited Apr. 26, 2025) (emphasis added); *see also* <https://github.com/FerretDB/FerretDB> (last visited Apr. 18, 2025).

¹² PGConf India, PGConf India 2024 - Moving MongoDB Workloads to PostgreSQL by Peter Farkas (FerretDB), Youtube (Mar. 22, 2024), <https://www.youtube.com/watch?v=UG5oXy0QuEM>.

developer community. They distort these individuals’ understanding of the companies’ comparative product offerings and therefore may cause them to use FerretDB’s products, and/or *not* use MongoDB’s products, based on false understandings—and to their own detriment. This is not a hypothetical scenario: Echoing FerretDB’s own misleading language, customers have described FerretDB (inaccurately) as a “drop-in replacement” for MongoDB,¹³ and as a “MongoDB alternative.”¹⁴ Unsurprisingly, therefore, MongoDB users—misled by FerretDB’s misrepresentations—report that they have sought to replace MongoDB with FerretDB but “failed.”¹⁵

36. FerretDB knows that its representations are false, misleading, and confusing to customers of MongoDB. FerretDB’s CEO Peter Farkas openly admits that his and FerretDB’s statements have caused “most users [to] expect that [FerretDB] is going to be the exact same thing, [but] unfortunately, that’s not the case. MongoDB has a lot of advanced features, a lot of features, which none of the alternatives implemented [in FerretDB].”¹⁶ Farkas’s concessions confirm that users “expect” the performance and features that FerretDB has prominently claimed to offer on its website and in other communications, but which it does not and cannot deliver.

¹³ Reddit (Dec. 8, 2021, 7:05 A.M.), https://www.reddit.com/r/programming/comments/rbtaul/ferretdb_a_truly_open_source_mongodb_alternative/ (last visited Apr. 23, 2025).

¹⁴ Reddit (May 15, 2024, 12:37 P.M.), https://www.reddit.com/r/PostgreSQL/comments/1cstowe/ferretdb_is_an_opensource_proxy_that_converts/ (last visited Apr. 23, 2025).

¹⁵ Reddit (Apr. 1, 2025, 7:03 A.M.), https://www.reddit.com/r/graylog/comments/1jowdgd/replace_mongodb_with_ferretdb/.

¹⁶ Contributor, *Oxygen Deprivation: FerretDB with Peter Farkas*, at 30:20 (Jan. 31, 2024), <https://www.contributor.fyi/ferretdb>.

FerretDB Has Harmed MongoDB's Trusted Brand, Including Its Trademarks

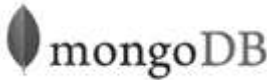
37. FerretDB is also causing undue harm to MongoDB's well-earned brand, including its trademarks. In the software community, a trusted brand is important because it serves as a signal of reliability, security, and long-term commitment to users and contributors alike. Free and source-available projects in particular often rely on voluntary adoption and collaboration. To that end, a strong, reputable brand reassures developers, businesses, and end users that the project will be maintained responsibly, that contributions will be managed fairly, and that the codebase can be depended on for critical applications. Uncertainty about and dilution of MongoDB's brand and associated trademarks stands to threaten user participation in its ecosystem and to harm customer adoption and retention.

38. **MongoDB Invests in Valuable Trademarks:** Since as early as 2008, MongoDB has adopted and used in commerce, including in Delaware, a family of MONGODB trademarks in connection with the MongoDB Database. As a result of MongoDB's extensive sales, advertising, and promotion of database services under its trademarks, the technology industry and consuming public have come to recognize MongoDB's trademarks as identifying a single source of high-quality products and services. Numerous media publications and competitors have recognized that MongoDB "makes the world's most popular general-purpose database."¹⁷ In 2014, an independent research firm named MongoDB as leader in the database industry,

¹⁷ Glenn Solomon, *Building An Enterprise Software Company*, Forbes (Feb. 25, 2021), <https://www.forbes.com/sites/glennsolomon/2021/02/25/building-an-enterprise-software-company-tips-for-founders-from-dev-ittycheria-ceo-of-mongodb/>; see also Jeffrey Erickson, *What Is MongoDB? An Expert Guide*, Oracle (Oct. 30, 2024), <https://www.oracle.com/database/mongodb/> ("MongoDB is a popular open source document database that's widely used in modern web and mobile applications.").

giving MongoDB the highest scores for market presence and recognizing MongoDB as “the most popular NoSQL document database.”¹⁸

39. MongoDB owns the following U.S. trademark registrations, among others, for use in connection with computer software for use in data management; database design and development; IT consulting services; technical support services; and educational services (collectively, the “MongoDB Marks”):

Mark	Reg. No.	Reg. Date	First Use in U.S. Commerce
MONGODB	3,930,955	March 15, 2011	August 12, 2008
MONGODB	4,523,518	April 29, 2014	February 19, 2010
MONGO	3,930,957	March 15, 2011	August 12, 2008
 mongoDB	3,930,956	March 15, 2011	August 12, 2008

40. A true and correct copy of the registration information for Reg. No. 3,930,955 is attached as **Exhibit A**.

41. A true and correct copy of the registration information for Reg. No. 4,523,518 is attached as **Exhibit B**.

42. A true and correct copy of the registration information for Reg. No. 3,930,957 is attached as **Exhibit C**.

43. A true and correct copy of the registration information for Reg. No. 3,930,956 is attached as **Exhibit D**.

¹⁸ *MongoDB Named a Leader by Independent Research Firm*, Yahoo Finance (Oct. 1, 2014), <https://finance.yahoo.com/news/mongodb-named-leader-independent-research-120000009.html>.

44. Pursuant to Section 7(b) of the Lanham Act, 15 U.S.C. § 1057(b), MongoDB's federal registration certificates are *prima facie* evidence of the validity of the above marks as well as MongoDB's ownership and exclusive right to use the MongoDB Marks in connection with the identified goods and services. All of the MongoDB Marks have achieved incontestable status under Section 15 of the Lanham Act, 15 U.S.C. § 1065.

45. The MongoDB Marks are inherently distinctive and consistently used in connection with the advertising, promotion, and sale of its goods and services. In advertising alone, MongoDB spent \$30.8 million in fiscal year 2025, \$29.7 million in fiscal year 2024, \$18.7 million in fiscal year 2023, \$18.0 million in fiscal year 2022, \$12.8 million in fiscal year 2021, \$7.6 million in fiscal year 2020, and \$5.1 million in fiscal year 2019.¹⁹

46. MongoDB targets customers and advertises the MongoDB Marks through print media, online ads, social media, blogs, whitepapers, case studies, webinars, newsletters, promotional emails, technology conferences, developer events, and partnerships and collaborations with other technology companies. MongoDB has also run television advertisements within the United States.

47. As a result of MongoDB's extensive advertising efforts, MongoDB has over 54,500 customers in more than 100 countries around the world. In fiscal year 2025 alone, MongoDB generated a total revenue of \$2.01 billion.

48. MongoDB exclusively uses and has invested significant resources to protect the MongoDB Marks. MongoDB's website lists the MongoDB Marks and provides a "Trademark

¹⁹ See MongoDB, Annual Report (Form 10-K) (Mar. 21, 2025); MongoDB, Annual Report (Form 10-K) (Mar. 18, 2022); MongoDB, Annual Report (Form 10-K) (Apr. 1, 2019).

Standards for Use” guide (“Trademark Standards”) for third parties.²⁰ MongoDB regularly monitors for improper third-party use of the MongoDB Marks and enforces its Trademark Standards against infringers and webhosts.

49. **FerretDB Misuses MongoDB’s Trademarks:** In an effort to take advantage of MongoDB’s brand and recognition, FerretDB has misused, muddled, and diluted the MongoDB Marks to promote FerretDB. FerretDB made clear its ill intent when it first launched in 2021, calling itself *MangoDB*. According to Defendants’ CEO and co-founder, Peter Farkas, “one of the reasons [they] chose” the company’s original name, ‘MangoDB’ is because, *obviously, MongoDB comes into mind.*”²¹ From its very beginning, FerretDB actively sought to trade on MongoDB’s goodwill.

50. FerretDB’s pattern of misusing MongoDB’s brand has continued unabated to the present day. For one, FerretDB characterizes its product as a “MongoDB alternative,” allowing users to “turn their Postgres services into MongoDB as a service.”²² FerretDB’s home page also displays “MONGODB” nearly as prominently as FerretDB’s logo. In fact, FerretDB uses the MongoDB Marks at least *nineteen times* on its home page alone.²³

51. For example, FerretDB claims on its homepage that it allows customers to “*use MongoDB drivers seamlessly with PostgreSQL,*” “*[l]everage the power of the MongoDB ecosystem,*” and “*[u]se all tools and drivers you would with MongoDB.*” With these

²⁰ MongoDB, Trademark Standards for Use, <https://www.mongodb.com/legal/trademark-usage-guidelines> (last visited Apr. 25, 2025).

²¹ Document Database Community, The Current State of MongoDB Alternatives, and 2Y of FerretDB, Youtube at 36:10 (Sept. 19, 2023), https://www.youtube.com/watch?v=43F1I9l6_gA.

²² PGConf India, PGConf India 2024 - Moving MongoDB Workloads to PostgreSQL by Peter Farkas (FerretDB), Youtube (Mar. 22, 2024), <https://www.youtube.com/watch?v=UG5oXy0QuEM>.

²³ <https://www.ferretdb.com/> (last visited Apr. 29, 2025).

statements, FerretDB is not merely identifying MongoDB, but suggesting to users that it *offers* MongoDB products and services.

52. FerretDB's conspicuous and repeated use of the MongoDB Marks in these ways, and for FerretDB's own gains, is likely to cause user confusion—especially given FerretDB's extensive false representations to users that its product is much the *same* as MongoDB's. Of course, it is FerretDB's *goal* to cause user confusion, as evidenced by its CEO's statement that it originally named the company MangoDB to cause users to think of MongoDB.

53. In fact, as discussed above, MongoDB is aware of evidence that some customers and users are in fact confused as a result of FerretDB's abuse of the MongoDB's Marks. To repeat: at least some users now believe FerretDB is a "MongoDB alternative" or a "drop-in replacement" for MongoDB.²⁴ Even FerretDB's CEO has admitted as much, stating that "most users expect that [MongoDB and FerretDB are] going to be the *exact same thing*."²⁵

54. Moreover, user confusion is exacerbated because the parties market competing goods through many of the same channels (*i.e.*, online and at technology events in the United States and all over the world).²⁶

²⁴ Reddit (May 15, 2024, 12:37 P.M.), https://www.reddit.com/r/PostgreSQL/comments/1cstowe/ferretdb_is_an_opensource_proxy_that_converts/ (last visited Apr. 23, 2025); Reddit (Dec. 8, 2021, 7:05 A.M.), https://www.reddit.com/r/programming/comments/rbtaul/ferretdb_a_truly_open_source_mongodb_alternative/ (last visited Apr. 23, 2025).

²⁵ Contributor, *Oxygen Deprivation: FerretDB with Peter Farkas*, at 30:19 (Jan. 31, 2024), <https://www.contributor.fyi/ferretdb>.

²⁶ See, e.g., @FerretDB, TechHub.social (Feb. 4, 2025, 3:34 AM) <https://techhub.social/@ferretdb/113945403597327665>; @FerretDB, TechHub.social (Feb. 1, 2025, 2:16 AM), <https://techhub.social/@ferretdb/113928107743240224>; @ferret_db, X (May 22, 2023, at 11:00 AM), https://x.com/ferret_db/status/1660707319523725312.

55. The results of FerretDB's intentional efforts to confuse consumers have even rippled through the media: An independent consumer publication has stated that "FerretDB can turn any Postgres database system into a MongoDB service provider."²⁷

56. In addition, FerretDB's repeated and ongoing attempts to associate its brand with MongoDB have reduced the distinctiveness of the MongoDB Marks. This is especially true given FerretDB's identical use of the MongoDB Marks, the inherent distinctiveness of the MongoDB Marks, MongoDB's otherwise exclusive use of the MongoDB Marks, and FerretDB's evident (and indeed, stated) improper intent in using the MongoDB Marks.

57. Moreover, FerretDB's conduct has damaged the positive reputation of the MongoDB Marks. FerretDB's inferior quality and attempts to associate itself with MongoDB have not gone unnoticed by FerretDB users, who have turned to the Internet to air their frustrations. FerretDB's use of MongoDB's trademarks has led to instances of actual confusion. Users misled by MongoDB have come to believe that FerretDB has the features, source attributes, and functionality of MongoDB. FerretDB's use of MongoDB's mark has also tarnished MongoDB's reputation. As a result of being associated with FerretDB's shoddy product, MongoDB's reputation and business have suffered.

58. Even if some of FerretDB's customers or potential customers eventually realize that FerretDB is not part of or related to MongoDB, the initial interest FerretDB is able to generate in its services from users who are drawn to FerretDB's website or other promotional or advertising materials because of its use of the MongoDB Marks is itself harmful to MongoDB,

²⁷ Joab Jackson, *FerretDB 2.0: Open Source MongoDB With PostgreSQL Power*, The Newstack (Mar. 5, 2025), <https://thenewstack.io/ferretdb-2-0-open-source-mongodb-alternative-with-postgresql-power/>.

FerretDB’s own users, and others in the developer community whose time is wasted and trust jeopardized.

FerretDB Ignores MongoDB’s Attempts to Engage, and Refuses to Stop Its Unfair and Unlawful Practices

59. On November 3, 2023, MongoDB sent FerretDB a letter asking it to stop its unlawful, harmful conduct. MongoDB informed FerretDB that, among other things, it was infringing MongoDB’s patents and trademarks and falsely describing itself and MongoDB.

60. After FerretDB failed to respond, MongoDB subsequently sent FerretDB a second letter on November 29, 2023. MongoDB again asked FerretDB to stop its unlawful and unfair conduct.

61. FerretDB continued to ignore MongoDB’s efforts to engage in good faith. FerretDB did not respond until April 16, 2024—nearly six months after MongoDB sent FerretDB its second letter.

62. Unfortunately, rather than negotiate in good faith with MongoDB, FerretDB has only escalated its illegal conduct. In 2025, FerretDB released version 2.0 of its product, which it touted as a “faster, more compatible MongoDB alternative,”²⁸ and the “result of a year-long collaboration” with Microsoft.²⁹ And FerretDB has also announced the imminent release of “FerretDB Cloud,” with which it claims users can “effortlessly deploy and scale FerretDB on

²⁸ Peter Farkas, *FerretDB Releases 2.0: A Faster, More Compatible MongoDB Alternative* (Jan. 23, 2025), <https://blog.ferretdb.io/ferretdb-releases-v2-faster-more-compatible-mongodb-alternative/>.

²⁹ Peter Farkas, *FerretDB 2.0 GA: Open Source MongoDB alternative, ready for production* (Mar. 5, 2025), <https://blog.ferretdb.io/ferretdb-v2-ga-open-source-mongodb-alternative-ready-for-production/>

AWS and GCP, with support for Microsoft Azure and additional cloud providers coming soon.”³⁰

63. Faced with FerretDB’s escalation, MongoDB sent a third and final letter, which included claim charts comparing FerretDB’s products to MongoDB’s asserted patent claims, on May 16, 2025. But FerretDB’s illegal conduct has continued unabated.

COUNT 1

PATENT INFRINGEMENT OF U.S. PATENT NO. 8,996,463

64. MongoDB hereby repeats and incorporates the allegations set forth above as though fully set forth herein.

65. MongoDB is the assignee and lawful owner of, and holds all right, title, and interest in the ’463 Patent, which issued March 31, 2015 and is entitled “Aggregation Framework System Architecture and Method.” A true and correct copy of the ’463 Patent is attached as **Exhibit E**.

66. The claims of the ’463 Patent are valid and enforceable.

67. The claims of the ’463 Patent are directed to patentable subject matter. The ’463 Patent is directed to an inventive aggregation framework system architecture and method to optimize aggregate operations over a database based on data dependencies. The ’463 Patent’s claimed aggregation framework system improves database functionality by, among other things, reducing at least some of the volume of data needed to complete an aggregation operation, reducing the data communicated between distributed portions of the database, reducing memory requirements for output/pipeline data, re-ordering execution of operations or commands within

³⁰ Peter Farkas, *FerretDB Releases 2.0: A Faster, More Compatible MongoDB Alternative* (Jan. 23, 2025), <https://blog.ferretdb.io/ferretdb-releases-v2-faster-more-compatible-mongodb-alternative/>.

the aggregation operation, enabling nested pipeline execution on array fields, and parallelizing aggregation operation execution. The claimed inventions provide specific concrete solutions to the problem of optimizing aggregation commands in non-traditional database architectures, which were previously cumbersome, inefficient, and led to additional complexity due to the structure of the underlying data in the database.

68. Claim 1 of the '463 Patent recites:

1. A computer implemented method for executing aggregation operations in a non-relational architecture, the method comprising the acts of
 - receiving, at database routing system, database requests from computer executed processes;
 - routing, by the database routing system, the database requests to a plurality of database access managers over a communication network;
 - identifying, by an aggregation engine, an aggregation operation including a plurality of data operations;
 - determining, by the aggregation engine, an optimization for execution of the aggregation operation;
 - modifying, by the aggregation engine, the plurality of data operations to optimize execution; and
 - communicating, by the aggregation engine, the aggregation operation having at least one optimization for execution.

'463 Patent, claim 1.

69. On information and belief, in violation of 35 U.S.C. § 271(a), FerretDB has directly infringed (both literally and under the doctrine of equivalents) and continues to directly infringe (both literally and under the doctrine of equivalents) one or more claims of the '463 Patent, including at least claim 1 of the '463 Patent, in the State of Delaware, in this District, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products and services that embody one or more of the

inventions claimed in the '463 Patent, including versions of FerretDB made, used, sold, or offered for sale in the past six years, which include but are not limited to those installed natively and those offered on the cloud (the "Accused Products"). As one example, FerretDB directly infringes through internal testing and development of the Accused Products, which implement elements that are identical or equivalent to each claimed element of the patented invention in claim 1 of the '463 Patent.

70. Each of the Accused Products implements a method recited in claim 1 of the '463 Patent. A claim chart comparing an independent claim of the '463 Patent to a representative Accused Product is attached as **Exhibit F**, which is hereby incorporated by reference in its entirety.

71. Further, on information and belief, FerretDB has actively induced and/or contributed to infringement of at least claim 1 of the '463 Patent in violation of at least 35 U.S.C. § 271(b) and (c).

72. Users of the Accused Products directly infringe at least claim 1 of the '463 Patent when they use the Accused Products in the ordinary, customary, and intended way.

73. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to consumers within the United States, and instructing and encouraging such customers to use the Accused Products in the ordinary, customary, and intended way, which FerretDB knew infringes at least claim 1 of the '463 Patent, or, alternatively, was willfully blind to the infringement.

74. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Products in the United States, which FerretDB knew infringes at least claim 1 of the '463 Patent, or, alternatively, was willfully blind to the infringement.

75. On information and belief, FerretDB actively advertises the Accused Products with instructions to users to encourage infringement.

76. For example, FerretDB describes the operation and implementation details of the Accused Products on its websites.³¹ On information and belief, FerretDB provides instructions to users to encourage infringement of the '463 Patent claims. For example, FerretDB provides instructions on its website to users on how to implement aggregation operations in an infringing manner.³²

77. As one example, FerretDB provides a description on its website³³ of how the Accused Products optimize aggregation operations:

³¹ See, e.g., <https://www.ferretdb.com/> (last visited Apr. 21, 2025); <https://blog.ferretdb.io/> (last visited Apr. 21, 2025); <https://docs.ferretdb.io/> (last visited Apr. 21, 2025).

³² <https://docs.ferretdb.io/usage/aggregations/> (last visited Apr. 21, 2025).

³³ <https://blog.ferretdb.io/ferretdb-fetches-data-query-pushdown/> (last visited Apr. 24, 2025).

Why does FerretDB need SQL query pushdowns?

As we aim to be as compatible with MongoDB drivers as possible, all operations, comparisons, data types, and commands need to be handled in the same fashion as MongoDB. Because of that, we cannot rely on SQL queries and filter data just in queries.

For example how would we compare values of the different types (considering BSON types comparison order)? The solution is to do all filtering operations on our own. That creates a need for fetching all the data from a storage layer, which for large collections, can be really inefficient and time-consuming.

Considering this fact, query pushdowns are a really important method for decreasing the data that we must fetch for every query. That's why FerretDB really can benefit from using them in such queries. Fortunately, we've managed to introduce the query pushdown with [this PR](#)!

78. On information and belief, the Accused Products have supported commands that implement aggregation operations in an infringing manner since at least the release of FerretDB 0.6.1.³⁴

79. On information and belief, in violation of 35 U.S.C. § 271(c), FerretDB's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least claim 1 of the '463 Patent, constituting a material part of the invention. On information and belief, FerretDB knows and has known the same to be especially made or especially adapted for use in an infringement of the '463 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use. For example, on information and belief, the Accused Products are not a staple article of commerce suitable for substantial non-infringing use, at least because they are especially designed and produced by FerretDB to understand and respond to aggregation

³⁴ <https://blog.ferretdb.io/new-release-ferretdb-0-6-1/> (last visited Apr. 26, 2025).

commands and queries in a manner claimed by the '463 Patent, and they are not capable of substantial non-infringing use.

80. FerretDB is not licensed or otherwise authorized to practice the claims of the '463 Patent.

81. Thus, by its acts, FerretDB has injured MongoDB and is liable to MongoDB for directly and/or indirectly infringing one or more claims of the '463 Patent, whether literally or under the doctrine of equivalents, including without limitation claim 1.

82. On information and belief, FerretDB commercializes and derives revenue and other monetary compensation from its infringement of one or more '463 Patent claims, including at least claim 1, by partnering with vendors, such as Civo, Tembo, Elestio, and Cozystack,³⁵ to use, distribute, and market the Accused Products, and by performing “comprehensive services for those who want to evaluate, run or customize FerretDB.” For example, FerretDB provides product support services to its customers and end users that include “everything from product evaluation, production down situations, 24/7 Q&A with your engineers, to complex technical troubleshooting.” FerretDB also provides product development services that “helps [its customers] create custom features into FerretDB, to ensure that it integrates with [the customers’] solutions.” FerretDB also provides “training and consulting” so that its customers can “quickly become[] an expert on running FerretDB.”³⁶

83. FerretDB has had knowledge of the '463 Patent and its infringement at least as early as November 3, 2023 and no later than the filing of this Complaint. This knowledge was compounded by subsequent letters sent on November 29, 2023, and May 16, 2025, laying out the

³⁵ <https://github.com/FerretDB/FerretDB?tab=readme-ov-file> (last visited Apr. 21, 2025).

³⁶ <https://www.ferretdb.com/services> (last visited Apr. 18, 2025).

precise manner in which the Accused Products infringed. FerretDB has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '463 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made FerretDB aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '463 Patent. On information and belief, discovery will reveal additional facts and circumstances from which FerretDB's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

84. Accordingly, FerretDB's infringement of the '463 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

85. As a result of FerretDB's infringement of the '463 Patent, MongoDB has suffered monetary damages and seeks recovery, in an amount to be proven at trial, adequate to compensate for FerretDB's infringement, but in no event less than a reasonable royalty with interest and costs.

86. On information and belief, FerretDB will continue to infringe the '463 Patent unless enjoined by this Court. FerretDB's infringement of MongoDB's rights under the '463 Patent will continue to damage MongoDB, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

COUNT 2

PATENT INFRINGEMENT OF U.S. PATENT NO. 9,262,462

87. MongoDB hereby repeats and incorporates the allegations set forth above as though fully set forth herein.

88. MongoDB is the assignee and lawful owner of, and holds all right, title, and interest in the '462 Patent, which issued February 16, 2016 and is entitled "Aggregation Framework System Architecture and Method." A true and correct copy of the '462 Patent is attached as **Exhibit G**.

89. The claims of the '462 Patent are valid and enforceable.

90. The claims of the '462 Patent are directed to patentable subject matter. Similar to the '463 Patent claims, the '462 Patent claims are directed to an inventive aggregation framework system architecture and method to optimize aggregate operations over a database based on data dependencies. The '462 Patent further optimizes aggregate operations by flattening data content prior to performing the claimed inventive aggregate operations. The '462 Patent's claimed aggregation framework system improves database functionality by, among other things, reducing at least some of the volume of data needed to complete an aggregation operation, reducing the data communicated between distributed portions of the database, reducing memory requirements for output/pipeline data, re-ordering execution of operations or commands within the aggregation operation, enabling nested pipeline execution on array fields, and parallelizing aggregation operation execution. The claimed inventions provide specific concrete solutions to the problem of optimizing aggregation commands in non-traditional database architectures, which were previously cumbersome, inefficient, and led to additional complexity due to the structure of the underlying data in the database.

91. Claim 1 of the '462 Patent recites:

1. A computer implemented method for consistent execution of aggregation expressions, in a non-relational architecture, the method comprising the acts of:

determining, by a computer system, an optimization for execution of an aggregation operation, wherein the aggregation operation includes a plurality of data operations on a non-relational database;

modifying, by the computer system, the plurality of data operations to optimize execution;

wherein determining the optimization for execution of the aggregation operation includes identifying an aggregation expression to execute as at least part of a data request on the non-relational database;

determining, by the computer system, a set of responsive data objects for input into the aggregation expression;

flattening, by the computer system, data content stored within each data object based on a specified attribute, wherein the act of flattening includes generating a new data object for each respective data record, stored within an array or as a reference to another data object within a respective data object, matching the specified attribute;

executing, by the computer system, the aggregation expression against the flattened data content.

'462 Patent, claim 1.

92. On information and belief, in violation of 35 U.S.C. § 271(a), FerretDB has directly infringed (both literally and under the doctrine of equivalents) and continues to directly infringe (both literally and under the doctrine of equivalents) one or more claims of the '462 Patent, including at least claim 1 of the '462 Patent, in the State of Delaware, in this District, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products and services that embody one or more of the inventions claimed in the '462 Patent, including the Accused Products. As one example, FerretDB directly infringes through internal testing and development of the Accused Products,

which implement elements that are identical or equivalent to each claimed element of the patented invention in claim 1 of the '462 Patent.

93. Each of the Accused Products implements a method recited in claim 1 of the '462 Patent. A claim chart comparing an independent claim of the '462 Patent to a representative Accused Product is attached as **Exhibit H**, which is hereby incorporated by reference in its entirety.

94. Further, on information and belief, FerretDB has actively induced and/or contributed to infringement of at least claim 1 of the '462 Patent in violation of at least 35 U.S.C. § 271(b) and (c).

95. Users of the Accused Products directly infringe at least claim 1 of the '462 Patent when they use the Accused Products in the ordinary, customary, and intended way.

96. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to consumers within the United States, and instructing and encouraging such customers to use the Accused Products in the ordinary, customary, and intended way, which FerretDB knew infringes at least claim 1 of the '462 Patent, or, alternatively, was willfully blind to the infringement.

97. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the Accused Products within the United States, by, directly or through intermediaries, instructing and

encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Products in the United States, which FerretDB knew infringes at least claim 1 of the '462 Patent, or, alternatively, was willfully blind to the infringement.

98. On information and belief, FerretDB actively advertises the Accused Products with instructions to users to encourage infringement.

99. For example, FerretDB describes the Accused Products on its websites.³⁷ On information and belief, FerretDB provides instructions to users to encourage infringement of the '462 Patent claims. For example, FerretDB provides instructions on its website to users on how to implement aggregation operations in an infringing manner.³⁸

100. On information and belief, the Accused Products provide several commands that operate to meet the optimization and flattening limitations of the '462 Patent claims. For example, the Accused Products support aggregation stages, such as \$group and \$unwind.³⁹

³⁷ See, e.g., <https://www.ferretdb.com/> (last visited Apr. 21, 2025); <https://blog.ferretdb.io/> (last visited Apr. 21, 2025); <https://docs.ferretdb.io/> (last visited Apr. 21, 2025).

³⁸ <https://docs.ferretdb.io/usage/aggregations/> (last visited Apr. 21, 2025).

³⁹ <https://docs.ferretdb.io/reference/aggregation-operations/aggregation-stages/> (last visited Apr. 24, 2025).

Version: v2.1

Aggregation stages

Aggregation stages are a series of one or more processes in a pipeline that acts upon the returned result of the previous stage, starting with the input documents. Some of the aggregation stages include:

Aggregation stages	Description
<code>\$count</code>	Returns the count of all matched documents in a specified query
<code>\$group</code>	Groups documents based on specific value or expression and returns a single document for each group
<code>\$limit</code>	Limits specific documents and passes the rest to the next stage
<code>\$match</code>	Acts as a <code>find</code> operation by only returning documents that match a specified query to the next stage
<code>\$project</code>	Specifies the fields in a document to pass to the next stage in the pipeline
<code>\$skip</code>	Skips a specified <code>n</code> number of documents and passes the rest to the next stage
<code>\$sort</code>	Sorts and returns all the documents based on a specified order
<code>\$unset</code>	Specifies the fields to be removed/excluded from a document
<code>\$unwind</code>	Deconstructs and returns a document for every element in an array field

On information and belief, the Accused Products have supported infringing stages, such as `$group` and `$unwind`, since at least the release of FerretDB v1.0.⁴⁰

101. On information and belief, in violation of 35 U.S.C. § 271(c), FerretDB's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least claim 1 of the '462 Patent, constituting a material part of the invention. On information and belief, FerretDB knows and has known the same to be especially made or especially adapted for use in an infringement of the '462 Patent, and such components are not a staple article or commodity of commerce suitable for substantial

⁴⁰ <https://blog.ferretdb.io/ferretdb-1-0-ga-opensource-mongodb-alternative/> (last visited Apr. 24, 2025).

noninfringing use. For example, on information and belief, the Accused Products are not a staple article of commerce suitable for substantial non-infringing use, at least because they are especially designed and produced by FerretDB to understand and respond to aggregation commands and queries in a manner claimed by the '462 Patent, and they are not capable of substantial non-infringing use.

102. FerretDB is not licensed or otherwise authorized to practice the claims of the '462 Patent.

103. Thus, by its acts, FerretDB has injured MongoDB and is liable to MongoDB for directly and/or indirectly infringing one or more claims of the '462 Patent, whether literally or under the doctrine of equivalents, including without limitation claim 1.

104. On information and belief, FerretDB commercializes and derives revenue and other monetary compensation from its infringement of one or more '462 Patent claims, including at least claim 1, by partnering with other vendors, such as Civo, Tembo, Elestio, and Cozystack,⁴¹ to use, distribute, and market the Accused Products, and by performing “comprehensive services for those who want to evaluate, run or customize FerretDB.” For example, FerretDB provides product support services to its customers and end users that include “everything from product evaluation, production down situations, 24/7 Q&A with your engineers, to complex technical troubleshooting.” FerretDB also provides product development services that “helps [its customers] create custom features into FerretDB, to ensure that it integrates with [the customers’] solutions.” FerretDB also provides “training and consulting” so that its customers can “quickly become[] an expert on running FerretDB.”⁴²

⁴¹ <https://github.com/FerretDB/FerretDB?tab=readme-ov-file> (last visited Apr. 21, 2025).

⁴² <https://www.ferretdb.com/services> (last visited Apr. 18, 2025).

105. FerretDB has had knowledge of the '462 Patent and its infringement at least as early as November 3, 2023 and no later than the filing of this Complaint. This knowledge was compounded by subsequent letters sent on November 29, 2023, and May 16, 2025, laying out the precise manner in which the Accused Products infringed. FerretDB has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '462 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made FerretDB aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '462 Patent. On information and belief, discovery will reveal additional facts and circumstances from which FerretDB's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

106. Accordingly, FerretDB's infringement of the '462 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

107. As a result of FerretDB's infringement of the '462 Patent, MongoDB has suffered monetary damages and seeks recovery, in an amount to be proven at trial, adequate to compensate for FerretDB's infringement, but in no event less than a reasonable royalty with interest and costs.

108. On information and belief, FerretDB will continue to infringe the '462 Patent unless enjoined by this Court. FerretDB's infringement of MongoDB's rights under the '462 Patent will continue to damage MongoDB, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

COUNT 3

PATENT INFRINGEMENT OF U.S. PATENT NO. 10,031,956

109. MongoDB hereby repeats and incorporates the allegations set forth above as though fully set forth herein.

110. MongoDB is the assignee and lawful owner of, and holds all right, title, and interest in the '956 Patent, which issued July 24, 2018 and is entitled "Aggregation Framework System Architecture and Method." A true and correct copy of the '956 Patent is attached as **Exhibit I.**

111. The claims of the '956 Patent are valid and enforceable.

112. The claims of the '956 Patent are directed to patentable subject matter. Similar to the '463 Patent claims, the '956 Patent claims are directed to an inventive aggregation framework system architecture and method to optimize aggregate operations over a database based on data dependencies. The '956 Patent further optimizes aggregate operations by flattening data content prior to performing the claimed inventive aggregate operations. The '956 Patent's claimed aggregation framework system improves database functionality by, among other things, reducing at least some of the volume of data needed to complete an aggregation operation, reducing the data communicated between distributed portions of the database, reducing memory requirements for output/pipeline data, re-ordering execution of operations or commands within the aggregation operation, enabling nested pipeline execution on array fields, and parallelizing aggregation operation execution. The claimed inventions provide specific concrete solutions to the problem of implementing aggregation commands in non-traditional database architectures, which were previously cumbersome, inefficient, and led to additional complexity due to the structure of the underlying data in the database.

113. Claim 1 of the '956 Patent recites:

1. A computer implemented method for consistent execution of aggregation expressions, the method comprising the acts of:

determining, by a computer system, an optimization for execution of an aggregation operation, wherein the aggregation operation includes a plurality of data operations on a non-relational database including an aggregation expression operation;

determining, by the computer system, a set of responsive data objects for input into the aggregation expression;

flattening, by the computer system, data content stored within each data object of the set of responsive data objects based on a specified attribute, wherein the act of flattening includes generating a new data object for each respective data record, stored within an array or as a reference to another data object within a respective data object, matching the specified attribute;

executing, by the computer system, the aggregation expression against the flattened data content.

'956 Patent, claim 1.

114. On information and belief, in violation of 35 U.S.C. § 271(a), FerretDB has directly infringed (both literally and under the doctrine of equivalents) and continues to directly infringe (both literally and under the doctrine of equivalents) one or more claims of the '956 Patent, including at least claim 1 of the '956 Patent, in the State of Delaware, in this District, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products and services that embody one or more of the inventions claimed in the '956 Patent, including the Accused Products. As one example, FerretDB directly infringes through internal testing and development of the Accused Products, which implement elements that are identical or equivalent to each claimed element of the patented invention in claim 1 of the '956 Patent.

115. Each of the Accused Products implements a method recited in claim 1 of the '956 Patent. A claim chart comparing an independent claim of the '956 Patent to a representative Accused Product is attached as **Exhibit J**, which is hereby incorporated by reference in its entirety.

116. Further, on information and belief, FerretDB has actively induced and/or contributed to infringement of at least claim 1 of the '956 Patent in violation of at least 35 U.S.C. § 271(b) and (c).

117. Users of the Accused Products directly infringe at least claim 1 of the '956 Patent when they use the Accused Products in the ordinary, customary, and intended way.

118. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to consumers within the United States, and instructing and encouraging such customers to use the Accused Products in the ordinary, customary, and intended way, which FerretDB knew infringes at least claim 1 of the '956 Patent, or, alternatively, was willfully blind to the infringement.

119. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the Accused Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Products in the United States, which FerretDB knew

infringes at least claim 1 of the '956 Patent, or, alternatively, was willfully blind to the infringement.

120. On information and belief, FerretDB actively advertises the Accused Products with instructions to users to encourage infringement.

121. For example, FerretDB describes the Accused Products on its websites.⁴³ On information and belief, FerretDB provides instructions to users to encourage infringement of the '956 Patent claims. For example, FerretDB provides instructions on its website to users on how to implement aggregation operations in an infringing manner.⁴⁴

122. On information and belief, the Accused Products provide several commands that operate to meet the optimization and flattening limitations of the '956 Patent claims. For example, the Accused Products support aggregation commands, such as \$group and \$unwind:⁴⁵

⁴³ See, e.g., <https://www.ferretdb.com/> (last visited Apr. 21, 2025); <https://blog.ferretdb.io/> (last visited Apr. 21, 2025); <https://docs.ferretdb.io/> (last visited Apr. 21, 2025).

⁴⁴ <https://docs.ferretdb.io/usage/aggregations/> (last visited Apr. 21, 2025).

⁴⁵ <https://docs.ferretdb.io/reference/aggregation-operations/aggregation-stages/> (last visited Apr. 24, 2025).

Version: v2.1

Aggregation stages

Aggregation stages are a series of one or more processes in a pipeline that acts upon the returned result of the previous stage, starting with the input documents. Some of the aggregation stages include:

Aggregation stages	Description
<code>\$count</code>	Returns the count of all matched documents in a specified query
<code>\$group</code>	Groups documents based on specific value or expression and returns a single document for each group
<code>\$limit</code>	Limits specific documents and passes the rest to the next stage
<code>\$match</code>	Acts as a <code>find</code> operation by only returning documents that match a specified query to the next stage
<code>\$project</code>	Specifies the fields in a document to pass to the next stage in the pipeline
<code>\$skip</code>	Skips a specified <code>n</code> number of documents and passes the rest to the next stage
<code>\$sort</code>	Sorts and returns all the documents based on a specified order
<code>\$unset</code>	Specifies the fields to be removed/excluded from a document
<code>\$unwind</code>	Deconstructs and returns a document for every element in an array field

On information and belief, the Accused Products have supported infringing commands, such as `$group` and `$unwind`, since at least the release of FerretDB v1.0.⁴⁶

123. On information and belief, in violation of 35 U.S.C. § 271(c), FerretDB's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least claim 1 of the '956 Patent, constituting a material part of the invention. On information and belief, FerretDB knows and has known the same to be especially made or especially adapted for use in an infringement of the '956 Patent, and such components are not a staple article or commodity of commerce suitable for substantial

⁴⁶ <https://blog.ferretdb.io/ferretdb-1-0-ga-opensource-mongodb-alternative/> (last visited Apr. 24, 2025).

noninfringing use. For example, on information and belief, the Accused Products are not a staple article of commerce suitable for substantial non-infringing use, at least because they are especially designed and produced by FerretDB to understand and respond to aggregation commands and queries in a manner claimed by the '956 Patent, and they are not capable of substantial non-infringing use.

124. FerretDB is not licensed or otherwise authorized to practice the claims of the '956 Patent.

125. Thus, by its acts, FerretDB has injured MongoDB and is liable to MongoDB for directly and/or indirectly infringing one or more claims of the '956 Patent, whether literally or under the doctrine of equivalents, including without limitation claim 1.

126. On information and belief, FerretDB commercializes and derives revenue and other monetary compensation from its infringement of one or more '956 Patent claims, including at least claim 1, by partnering with other vendors, such as Civo, Tembo, Elestio, and Cozystack,⁴⁷ to use, distribute, and market the Accused Products, and by performing “comprehensive services for those who want to evaluate, run or customize FerretDB.” For example, FerretDB provides product support services to its customers and end users that include “everything from product evaluation, production down situations, 24/7 Q&A with your engineers, to complex technical troubleshooting.” FerretDB also provides product development services that “helps [its customers] create custom features into FerretDB, to ensure that it integrates with [the customers’] solutions.” FerretDB also provides “training and consulting” so that its customers can “quickly become[] an expert on running FerretDB.”⁴⁸

⁴⁷ <https://github.com/FerretDB/FerretDB?tab=readme-ov-file> (last visited Apr. 21, 2025).

⁴⁸ <https://www.ferretdb.com/services> (last visited Apr. 18, 2025).

127. FerretDB has had knowledge of the '956 Patent and its infringement at least as early as November 3, 2023 and no later than the filing of this Complaint. This knowledge was compounded by subsequent letters sent on November 29, 2023, and May 16, 2025, laying out the precise manner in which the Accused Products infringed. FerretDB has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or, alternatively, through its willfully blind disregard of the '956 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made FerretDB aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '956 Patent. On information and belief, discovery will reveal additional facts and circumstances from which FerretDB's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

128. Accordingly, FerretDB's infringement of the '956 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

129. As a result of FerretDB's infringement of the '956 Patent, MongoDB has suffered monetary damages and seeks recovery, in an amount to be proven at trial, adequate to compensate for FerretDB's infringement, but in no event less than a reasonable royalty with interest and costs.

130. On information and belief, FerretDB will continue to infringe the '956 Patent unless enjoined by this Court. FerretDB's infringement of MongoDB's rights under the '956 Patent will continue to damage MongoDB, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

COUNT 4

PATENT INFRINGEMENT OF U.S. PATENT NO. 10,866,868

131. MongoDB hereby repeats and incorporates the allegations set forth above as though fully set forth herein.

132. MongoDB is the assignee and lawful owner of, and holds all right, title, and interest in the '868 Patent, which issued December 15, 2020 and is entitled "Systems and Methods for Optimization of Database Operations." A true and correct copy of the '868 Patent is attached as **Exhibit K**.

133. The claims of the '868 Patent are valid and enforceable.

134. The claims of the '868 Patent are directed to patentable subject matter. The '868 Patent is directed to an inventive distributed database system and method where it can determine if a write operation failed as a result of an error that occurred during that operation and then re-execute that same write operation in that case. This improves database functionality by, among other things, increasing the reliability of a write operation in a distributed system and minimizing disruptions and overhead that are incurred as a result of failed writes. This in turn allows database administrators more freedom to conduct operations such as system upgrades, which cause nodes within the distributed database to become temporarily unavailable and, without retryable writes, would cause a significant number of write failures. The claimed inventions provide specific concrete solutions to the problem of failed write operations because they involve the use of a specific architecture—a replica set consisting of a primary node and at least one secondary node—as well an operational log, logic which identifies write failures by determining the occurrence of an error during the write operation, and logic which triggers re-execution of the write command in response.

135. Claim 12 of the '868 Patent recites:

12. A computer-implemented method of managing a database, the method comprising acts of:

storing data in a distributed database having a dynamic schema architecture, the storing comprising storing a replica set hosting a respective shard of data;

performing, by a primary node of the replica set, write operations on the distributed database wherein the primary node includes an operational log;

replicating, by at least one secondary node of the replica set, write operations performed by the primary node, comprising retrieving the operational log and replicating write operations in the operational log performed by the primary node;

receiving, by at least one processor from a client system, a submission of a write operation to perform on the distributed database;

executing, by the at least one processor, the submitted write operation at least in part by transmitting a command to the primary node to perform the write operation;

determining, by the at least one processor, that the execution of the write operation failed responsive to determining occurrence of an error during execution of the write operation; and

triggering, by the at least one processor, re-execution of the submitted write operation responsive to determining that the execution of the write operation failed at least in part by re-transmitting the command to the primary node to perform the write operation.

'868 Patent, claim 12.

136. On information and belief, in violation of 35 U.S.C. § 271(a), FerretDB has directly infringed (both literally and under the doctrine of equivalents) and continues to directly infringe (both literally and under the doctrine of equivalents) one or more claims of the '868 Patent, including at least claim 12 of the '868 Patent, in the State of Delaware, in this District, and elsewhere in the United States by, among other things, making, using, selling, offering for sale, and/or importing into the United States products and services that embody one or more of the inventions claimed in the '868 Patent, including the Accused Products. As one example,

FerretDB directly infringes through internal testing and development of the Accused Products, which implement elements that are identical or equivalent to each claimed element of the patented invention in claim 12 of the '868 Patent.

137. Each of the Accused Products implements a method recited in claim 12 of the '868 Patent. A claim chart comparing an independent claim of the '956 Patent to a representative Accused Product is attached as **Exhibit L**, which is hereby incorporated by reference in its entirety.

138. Further, on information and belief, FerretDB has actively induced and/or contributed to infringement of at least claim 12 of the '868 Patent in violation of at least 35 U.S.C. § 271(b) and (c).

139. Users of the Accused Products directly infringe at least claim 12 of the '868 Patent when they use the Accused Products in the ordinary, customary, and intended way.

140. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) include, without limitation and with specific intent to encourage infringement, knowingly inducing consumers to use the Accused Products within the United States in the ordinary, customary, and intended way by, directly or through intermediaries, supplying the Accused Products to consumers within the United States, and instructing and encouraging such customers to use the Accused Products in the ordinary, customary, and intended way, which FerretDB knew infringes at least claim 12 of the '868 Patent, or, alternatively, was willfully blind to the infringement.

141. On information and belief, FerretDB's inducements in violation of 35 U.S.C. § 271(b) further include, without limitation and with specific intent to encourage the infringement, knowingly inducing customers to commit acts of infringement with respect to the Accused

Products within the United States, by, directly or through intermediaries, instructing and encouraging such customers to import, make, use, sell, offer to sell, or otherwise commit acts of infringement with respect to the Accused Products in the United States, which FerretDB knew infringes at least claim 12 of the '868 Patent, or, alternatively, was willfully blind to the infringement.

142. On information and belief, FerretDB actively advertises the Accused Products with instructions to users to encourage infringement.

143. For example, FerretDB describes the Accused Products on its websites.⁴⁹ On information and belief, FerretDB provides instructions to users to encourage infringement of the '868 Patent claims.

144. On information and belief, in violation of 35 U.S.C. § 271(c), FerretDB's contributory infringement further includes offering to sell or selling within the United States, or importing into the United States, components of the patented invention of and/or a material or apparatus for use in practicing at least claim 1 of the '868 Patent, constituting a material part of the invention. On information and belief, FerretDB knows and has known the same to be especially made or especially adapted for use in an infringement of the '868 Patent, and such components are not a staple article or commodity of commerce suitable for substantial noninfringing use. For example, on information and belief, the Accused Products are not a staple article of commerce suitable for substantial non-infringing use, at least because they are especially designed and produced by FerretDB to understand and respond to write failures in a manner claimed by the '868 Patent, and they are not capable of substantial non-infringing use.

⁴⁹ See, e.g., <https://www.ferretdb.com/> (last visited Apr. 21, 2025); <https://blog.ferretdb.io/> (last visited Apr. 21, 2025); <https://docs.ferretdb.io/> (last visited Apr. 21, 2025).

145. FerretDB is not licensed or otherwise authorized to practice the claims of the '868 Patent.

146. Thus, by its acts, FerretDB has injured MongoDB and is liable to MongoDB for directly and/or indirectly infringing one or more claims of the '868 Patent, whether literally or under the doctrine of equivalents, including without limitation claim 1.

147. On information and belief, FerretDB commercializes and derives revenue and other monetary compensation from its infringement of one or more '868 Patent claims, including at least claim 1, by partnering with other vendors, such as Civo, Tembo, Elestio, and Cozystack,⁵⁰ to use, distribute, and market the Accused Products, and by performing “comprehensive services for those who want to evaluate, run or customize FerretDB.” For example, FerretDB provides product support services to its customers and end users that include “everything from product evaluation, production down situations, 24/7 Q&A with your engineers, to complex technical troubleshooting.” FerretDB also provides product development services that “helps [its customers] create custom features into FerretDB, to ensure that it integrates with [the customers’] solutions.” FerretDB also provides “training and consulting” so that its customers can “quickly become[] an expert on running FerretDB.”⁵¹

148. FerretDB has had knowledge of the '868 Patent and its infringement at least as early as November 3, 2023 and no later than the filing of this Complaint. This knowledge was compounded by subsequent letters sent on November 29, 2023, and May 16, 2025, laying out the precise manner in which the Accused Products infringed. FerretDB has had, and continues to have, the specific intent to infringe, through its deliberate and intentional infringement or,

⁵⁰ <https://github.com/FerretDB/FerretDB?tab=readme-ov-file> (last visited April 21, 2025).

⁵¹ <https://www.ferretdb.com/services> (last visited Apr. 18, 2025).

alternatively, through its willfully blind disregard of the '868 Patent by knowing there was a high probability of infringement but taking deliberate actions to avoid confirming that infringement. The filing of this action has also made FerretDB aware of the unjustifiably high risk that its actions constituted and continue to constitute infringement of the '868 Patent. On information and belief, discovery will reveal additional facts and circumstances from which FerretDB's knowledge and intent to infringe (or willful indifference), both before and after the filing of this action, may be inferred.

149. Accordingly, FerretDB's infringement of the '868 Patent has been and continues to be deliberate, intentional, and willful, and this is therefore an exceptional case warranting an award of enhanced damages and attorneys' fees and costs pursuant to 35 U.S.C. §§ 284 and 285.

150. As a result of FerretDB's infringement of the '868 Patent, MongoDB has suffered monetary damages and seeks recovery, in an amount to be proven at trial, adequate to compensate for FerretDB's infringement, but in no event less than a reasonable royalty with interest and costs.

151. On information and belief, FerretDB will continue to infringe the '868 Patent unless enjoined by this Court. FerretDB's infringement of MongoDB's rights under the '868 Patent will continue to damage MongoDB, causing irreparable harm for which there is no adequate remedy at law, unless enjoined by this Court.

COUNT 5

FALSE ADVERTISING UNDER THE LANHAM ACT (15 U.S.C. § 1125(a))

152. MongoDB hereby repeats and incorporates the allegations set forth above as though fully set forth herein.

153. FerretDB has made false or misleading statements in its commercial advertisements that promote FerretDB's offering and disparage MongoDB's products, including but not limited to its MongoDB database.

154. FerretDB's false and misleading advertisements are material and likely to influence customers' purchasing decisions about FerretDB's service and products, thereby diverting revenues from MongoDB. FerretDB's false and misleading advertisements are also likely to lessen the goodwill associated with MongoDB by associating the MongoDB Marks with the inferior products offered by FerretDB.

155. FerretDB's wrongful activities are causing MongoDB immediate irreparable injury. Upon information and belief, FerretDB will continue its wrongful activities to the continued and irreparable injury of MongoDB unless enjoined by this Court. MongoDB has no adequate remedy at law. MongoDB is therefore entitled to a preliminary and permanent injunction pursuant to 15 U.S.C. § 1116.

156. Pursuant to 15 U.S.C. § 1117, MongoDB is also entitled to recover FerretDB's profits, MongoDB's ascertainable damages, and MongoDB's costs of suit. FerretDB's willful use of the MongoDB Marks without excuse or justification renders this an exceptional case and entitles MongoDB to its reasonable attorneys' fees.

COUNT 6

TRADEMARK DILUTION UNDER DELAWARE LAW (6 DEL. C. § 3313 ET SEQ.)

157. MongoDB hereby repeats and incorporates the allegations set forth above as though fully set forth herein.

158. MongoDB owns valid rights in the MongoDB Marks.

159. FerretDB's unauthorized use of the MongoDB Marks is likely to injure MongoDB's business reputation and has diluted, or is likely to dilute, the distinctive quality of the MongoDB Marks in violation of the Delaware Trademark Act, 6 Del. C. § 3313.

160. FerretDB intended to trade on MongoDB's reputation and to dilute the MongoDB Marks, or acted with reason to know or was willfully blind as to the consequences of its actions.

161. As a result of FerretDB's conduct, MongoDB has suffered and continues to suffer damages in an amount yet to be determined.

162. Upon information and belief, the acts of dilution by FerretDB have resulted, and are currently resulting, in substantial unjust profits and unjust enrichment for FerretDB in an amount yet to be determined.

163. FerretDB's wrongful acts have caused and will continue to cause MongoDB irreparable harm. MongoDB has no adequate remedy at law for FerretDB's violations.

164. MongoDB is entitled to a judgment enjoining and restraining FerretDB from engaging in further acts of injury to MongoDB's business reputation or dilution.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff MongoDB prays for relief, declaration, and judgment that:

- a. The Asserted Patents are valid and enforceable;
- b. FerretDB has infringed of one or more claims of the Asserted Patents;
- c. FerretDB's infringement of the Asserted Patents was willful;
- d. MongoDB is entitled to all appropriate damages under 35 U.S.C. § 284 for FerretDB's infringement of the Asserted Patents, including pre-judgment and post-judgment interest, costs, and disbursements as justified under 35 U.S.C. § 284 and, if necessary to

adequately compensate MongoDB for FerretDB's infringement, an accounting, including an award of treble damages by reason of FerretDB's willful infringement of the Asserted Patents;

e. MongoDB is entitled to preliminary and permanent injunctive relief enjoining FerretDB from continuing to infringe the Asserted Patents;

f. FerretDB has willfully engaged in false advertising in violation of the Lanham Act;

g. MongoDB is entitled to all appropriate remedies under 15 U.S.C. § 1117 for FerretDB's false advertising, including an award of FerretDB's profits, any damages sustained by MongoDB, corrective advertising, costs, and treble damages;

h. MongoDB is entitled to preliminary and permanent injunctive relief enjoining FerretDB from false advertising;

i. FerretDB's unauthorized use of the MongoDB Marks, as alleged in this Complaint, constitutes trademark dilution in violation of Delaware law;

j. MongoDB is entitled to an award of damages, including disgorgement, unjust enrichment, and unjust profits, in an amount to be determined at trial;

k. MongoDB is entitled to preliminary and permanent injunctive relief enjoining FerretDB from continuing to dilute the MongoDB Marks;

l. This case is exceptional under, *inter alia*, 35 U.S.C. § 285 and 15 U.S.C. § 1117 such that MongoDB is awarded its attorneys' fees, costs, and expenses incurred in this action; and

m. MongoDB is entitled to any and all other legal and equitable relief as may be available under law and that the Court may deem proper.

DEMAND FOR JURY TRIAL

Plaintiff MongoDB, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

Respectfully submitted,

/s/ Nathan R. Hoeschen

Nathan R. Hoeschen (No. 6232)

Virginia K. Lynch (No. 7423)

SHAW KELLER LLP

I.M. Pei Building

1105 North Market Street, 12th Floor

Wilmington, DE 19801

(302) 298-0700

nhoeschen@shawkeller.com

glynch@shawkeller.com

Attorneys for Plaintiff MongoDB, Inc.

OF COUNSEL:

David Silbert

Sharif E. Jacob

Kristen E. Lovin

Ryan J. Hayward

Theresa M. Dawson

Yegina Whang

KEKER, VAN NEST & PETERS LLP

633 Battery Street

San Francisco, CA 94111-1809

(415) 391-5400

Dated: May 23, 2025